Probabilistic Diagnosability of Hybrid Systems

Yi Deng*
Department of Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute

A. Agung Julius
Department of Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute

Alessandro D'Innocenzo
Department of Engineering and Information Sciences and Mathematics
Center of Excellence DEWS
University of L'Aquila, Italy

ABSTRACT

The model-based fault diagnosability analysis is concerned with the timely detection and isolation of faults by using the system model and observations of the system output. In this paper, we propose the $(\delta_d, \delta_m, \alpha)$-diagnosability notion for hybrid systems with probabilistic reset, where the faults are diagnosed by observing the timed event sequences. We also present an approach for the analysis of such diagnosability.

The $(\delta_d, \delta_m, \alpha)$-diagnosability notion characterizes the worst-case probability α of detecting and isolating faults within the maximum delay δ_d since their first occurrence, given the measurement uncertainty δ_m in observing the time intervals between observed events. We present a method of system abstraction, and prove a quantitative relation between the $(\delta_d, \delta_m, \alpha)$-diagnosability of the original system and the abstraction. The abstraction has only finitely many trajectories that extend to the end of the time horizon of interest, which allows us to practically calculate the diagnosability and construct the diagnoser.

1. INTRODUCTION

When a complex system operates, a fault may occur in any of its component. Detection and isolation of faults as quick as possible could keep the system from incurring severe damages, and even saves human lives. Fault detection and isolation comprise the major task in the fault diagnosis [9], which can be performed by comparing available measurements with the model information. Based on the system model, a fault has some particular pattern of anticipated measurements as its symptoms, called the *fault signature* [3]. Intuitively, fault diagnosability is determined by the discriminability of fault signatures. If the symptom of a fault is confounded with the normal system behavior, obviously fault detection cannot be made. Similarly, fault isolability requires the discriminability of different faults.

*YD and AAJ would like to acknowledge the support of NSF CAREER grant CNS-0953976.

HSCC '15 April 14 - 16, 2015, Seattle, WA, USA

http://dx.doi.org/10.1145/2728606.2728621

In this paper, we propose $(\delta_d, \delta_m, \alpha)$-diagnosability notion for hybrid systems, and present a methodology for analyzing such diagnosability. Since the discriminability of observed symptoms can be impaired by low accuracy of the actual measurements, the *measurement uncertainty* parameter δ_m is incorporated into the diagnosability notion. In addition, a delay parameter δ_d is specified, which represents the maximum delay for diagnosing the fault since its first occurrence. Relevant works can be found in [6, 8, 16].

If the system model involves probabilistic dynamics, then it could be the case that the symptoms are not completely discriminable, but the probability of misclassifying a symptom can be described. We thus propose the notion of $(\delta_d, \delta_m, \alpha)$-diagnosability as a weaker alternative to the logical diagnosability [13]. The $(\delta_d, \delta_m, \alpha)$-diagnosability notion requires that the probability of misclassification do not exceed the threshold $(1 - \alpha)$, while the logical diagnosability amounts to $\alpha = 1$.

There are some earlier works on diagnosability analysis of stochastic discrete event systems, for example [14, 15]. Our work is different from the probabilistic diagnosability of [14, 15] in the following way: [14, 15] are concerned with steady-state behaviors of stochastic discrete event systems, while the present paper investigates the finite-horizon fault diagnosability for hybrid systems with probabilistic reset. However, our notion of probabilistic diagnosability is inspired by the similar definition therein.

The present research is motivated by problem of network congestion diagnosis. The data communication network can be modeled as a hybrid system [2], where package loss events have probabilistic reset. By observing the timed events of package loss at only a few of the routers, we want to diagnose the congestion condition of the whole network. In Section 3.5, we illustrate this application with a simplified model.

Thus, the fault diagnosability analysis problem in this paper can be formulated as follows: Given a hybrid system with probabilistic reset, suppose one can only observe events with their timing. Moreover, only a subset of the events are modeled as observable, and the timing has limited measurement accuracy. We want to analyze whether the occurrence of any faulty event can be deduced within a limited time.

In the literature, there are some works on designing effective state estimators based on the measurement of contin-

uous systems states [12, 17], that is, to efficiently improve the state observability. In contrast, our approach focus on obtaining the knowledge of temporal discrete-event behavior of the hybrid system, given the limited observability. This makes it possible to diagnose faults by using the timing and order of *observable events* during the system operation. As it will be more apparent in later sections, such diagnosis by the observation of timed event sequences would require a pairwise comparison of all the possible sequences generated by the system. Due to the intrinsic complexity of hybrid dynamics, directly analyzing the original model is hard or even impossible. Therefore, an indirect approach by system abstraction is presented: In Section 2, we present a method to construct the system abstraction. Our system abstraction has finitely many trajectories extending to the time horizon of interest, so its probabilistic diagnosability can be easily calculated. In Section 3, we prove a quantitative relation between the probabilistic diagnosability of the hybrid system and the abstraction, which allows us to derive the former from the latter, and also build a diagnoser.

2. HYBRID SYSTEMS ABSTRACTION

2.1 Hybrid Systems Definition

In this paper we model hybrid systems basically in the same way as in [1] except that an event may have multiple target discrete states. Whenever an event is triggered, the discrete state is reset to one of the candidates with some probability. See Def. 1 below.

DEFINITION 1. *A hybrid autonomous system with probabilistic reset is a tuple* $H = (L \times X, L^0 \times X^0, D, E, Inv)$ *that consists of:*

- *A set* $L \times X$ *of hybrid states* (ℓ, x)*, where* $\ell \in L$ *is the discrete state, and* $x \in X$ *is the continuous state. Discrete states are also called locations.*
- *A set* $L^0 \times X^0 \subset L \times X$ *of initial states.*
- *D associates with each location* $\ell \in L$ *the autonomous continuous time-invariant dynamics,* $D_\ell : \dot{x} = D_\ell(x)$*. This differential equation is assumed to admit a unique global solution* $\xi_\ell(t, x_\ell^0)$*, where* ξ_ℓ *satisfies* $\frac{\partial \xi_\ell(t,x_\ell^0)}{\partial t} = D_\ell(\xi_\ell(t, x_\ell^0))$*, and* $\xi_\ell(0, x_\ell^0) = x_\ell^0$ *is the initial condition in* ℓ*.*
- $Inv : L \to X$ *associates with each location an invariant set* $Inv(\ell) \subset X$*. Only if the continuous state satisfies* $x \in Inv(\ell)$*, can the discrete state be at the location* ℓ*.*
- *E is a set of events. In each location* ℓ*, the system state evolves continuously according to* D_ℓ *until an event* $e := (\ell, [\ell'], g, r, p), e \in E$ *occurs. The event is guarded by* $g \in Inv(\ell)$*. Namely, a necessary condition for the occurrence of e is* $x \in g$*. Let* (ℓ, x) *denote the system state that triggers e. After the event, the location is reset to one of the possible targets,* $\ell' \in [\ell'] \subset L$*, and the continuous state is reset to* $r(\ell', x) \in Inv(\ell')$*. The probability of resetting the location to* ℓ' *is given by* $p(\ell', x)$*, where* $p : [\ell'] \times g \to (0, 1]$ *satisfies that for any fixed* $x \in g$*,* $\sum_{\ell' \in [\ell']} p(\ell', x) = 1$*.*

Let G_ℓ denotes the set of guards such that the associated events all have ℓ as the source location. Let $\partial Inv(\ell)_{out}$ denote part of the boundary $\partial Inv(\ell)$ where the continuous state is evolving outward $Inv(\ell)$, i.e., given $\xi_\ell(\tau, x_\ell^0) \in \partial Inv(\ell)_{out}$, for any $t > 0$, there exists $t_1 \in (0, t)$ such that $\xi_\ell(\tau + t_1, x_\ell^0) \notin Inv(\ell)$. We adopt the following assumptions:

1. Non-deadlocking. We require $\partial Inv(\ell)_{out} \subset G_\ell$ for all $\ell \in L$ in order to avoid deadlocking. Namely, whenever the continuous state is evolving outside $Inv(\ell)$, an event must be specified.
2. The initial set $L^0 \times X^0$ is compact. The initial state can vary in $L^0 \times X^0$ with non-determinism.
3. We assume that for any $\ell \in L$, $G_\ell \subset \partial Inv(\ell)_{out}$ holds; and for any $g_1, g_2 \in G_\ell$, g_1, g_2 are disjoint. With this assumption, the occurrence of events is deterministic: Whenever a guard is reached by the continuous state, an unique event is forced to occur. The present work can be extended to remove this assumption as discussed later.
4. Well-posedness. The differential equation $\dot{x} = D_\ell(x)$ admits a unique solution, namely, it satisfies the Lipschitz condition.
5. For any $e = (\ell, [\ell'], g, r, p) \in E$, for any fixed $\ell' \in [\ell']$, $p(\ell', x)$ is Lipschitz continuous with respect to x.
6. The system does not have Zeno behavior [10].
7. All the reset maps are continuous functions.

When the hybrid system runs, a sequence of events can be triggered. Only some of the events are observable, which are associated with observable output symbols $\psi \in \Psi_o$. Events that could not be observed are associated with the empty output symbol $\psi = \emptyset$. For convenience, we also define an initialization event $e^0 \notin E$ associated with the special output symbol ι (starting signal). Then a trajectory of the hybrid system can be defined as a sequence:

DEFINITION 2. *Given* $H = (L \times X, L^0 \times X^0, D, E, Inv)$*, a trajectory of H is*

$$\rho = (e^0, \ell^0, x^0, \tau^0), (e^1, \ell^1, x^1, \tau^1) \cdots = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^N,$$

such that

- $\forall i \geq 0$*,* $(\ell^i, x^i) \in L \times X$*, and* $(\ell^0, x^0) \in L^0 \times X^0$*;*
- $\forall i \geq 0$*,* $\tau^i \in \mathbb{R}_{\geq 0}$*, and* $\forall t \in [0, \tau^i]$*,* $\xi_{\ell^i}(t, x^i) \in Inv(\ell^i)$*;*
- $\forall i \geq 1$*,* $e^i = (\ell^{i-1}, [\ell^i], g^i, r^i, p^i) \in E$*,* $\ell^i \in [\ell^i]$*, and* $\xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1}) \in g^i$*,* $x^i = r^i(\ell^i, \xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1}))$*, i.e.,* (ℓ^i, x^i) *is the reset state.*

Suppose a trajectory $\rho' = \{(e'^i, \ell'^i, x'^i, \tau'^i)\}_{i=0}^{N'}$ is exactly the same as ρ except that ρ' lasts for a shorter time horizon, then we call ρ' a *sub-trajectory* of ρ. Formally, it should be satisfied $(e'^i, \ell'^i, x'^i, \tau'^i) = (e^i, \ell^i, x^i, \tau^i)$ for all $i \in [0, N'-1]$, and $(e'^{N'}, \ell'^{N'}, x'^{N'}) = (e^{N'}, \ell^{N'}, x^{N'})$; it is also required $N' < N, \tau'^{N'} \leq \tau^{N'}$, or $N' = N, \tau'^{N'} < \tau^{N'}$. In the special case $N' < N, \tau'^{N'} = \tau^{N'}$, ρ' is a prefix of ρ.

DEFINITION 3. *If a nonempty set of trajectories $[\rho]$ initiated from (ℓ^0, x^0) satisfies the following conditions, then it is called a trajectory tree of (ℓ^0, x^0).*

- $\forall \rho \in [\rho]$, *ρ is not a sub-trajectory of any $\rho' \in [\rho]$.*
- $\forall \rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N} \in [\rho]$ *with $N \geq 1$, for any $M \in \{1, \ldots, N\}$, there exists a subset of trajectories $[\rho]' \subset [\rho]$, such that every $\rho' = \{(e'^i, \ell'^i, x'^i, \tau'^i)\}_{i=0}^{N'} \in [\rho]'$ satisfies $N' \geq M$, and for all $i \in \{0, \ldots, M-1\}$, $(e'^i, \ell'^i, x'^i, \tau'^i) = (e^i, \ell^i, x^i, \tau^i)$, $e'^M = e^M$, moreover, the set $\{\ell'^M | \rho' = \{(e'^i, \ell'^i, x'^i, \tau'^i)\}_{i=0}^{N'} \in [\rho]'\}$ includes all the possible target locations of e^M.*

  In words, for any event e ever triggered by some $\rho \in [\rho]$, there must be a bunch of trajectories in $[\rho]$ that keep being the same as ρ until e is triggered, and then reset to every possible target location of e.

When we start a system run from (ℓ^0, x^0), as time horizon of the run indefinitely grows, if the outcome keeps being a sub-trajectory of certain trajectory ρ until the end of ρ, then we say the run is pinned down on ρ. The following proposition says that given a trajectory tree $[\rho]$, a system run is pinned down on one and only one ρ in $[\rho]$.

PROPOSITION 1. *Given a trajectory tree $[\rho]$ of (ℓ^0, x^0), every system run initiated from (ℓ^0, x^0) for sufficiently long time horizon results in a unique $\rho \in [\rho]$.*

PROOF. *Suppose a run from (ℓ^0, x^0) has been pinned down on both $\rho, \rho' \in [\rho]$ successively, then by definition ρ is a sub-trajectory of ρ'. This contradicts the definition of trajectory trees. So a run for sufficiently long time horizon can result in at most one trajectory in the tree.*

For all ℓ, $\dot{x} = D_\ell(x)$ admits a unique solution (Assumption 4); whenever a guard is reached, a unique event is forced to occur (Assumption 3); after the event, for any possible target location ℓ' of the event, there must be a trajectory $\rho \in [\rho]$ that takes ℓ' as its target location (Def. 3). Therefore, for any system run initiated from (ℓ^0, x^0), the outcome must match (keep being a sub-trajectory of) at least one $\rho \in [\rho]$ until the maximum time horizon of ρ. □

If Assumption 3 is eliminated, then a continuous state on a guard may or may not trigger an event, and the event that can be triggered by a continuous state is not unique. In that case, Prop. 1 does not hold anymore, but the results can be extended by defining a set of trajectory trees for (ℓ^0, x^0). For simplicity, we stick to determinism of events.

For each $(\ell^0, x^0) \in L^0 \times X^0$, given a trajectory tree of it, we define a sample space, whose elements are the trajectories that form the tree. These trajectories are called the *paths* of the tree. Each experiment is a system run from (ℓ^0, x^0) for sufficiently long time horizon, which results in one path by Prop. 1.

DEFINITION 4. *Given a trajectory tree $[\rho]$ of (ℓ^0, x^0), and a path $\rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N} \in [\rho]$, define the probability mass associated with ρ:*

$$P(\rho) = \prod_{i=1}^{N} p^i(\ell^i, \xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1})). \tag{1}$$

When $N = 0$ ($[\rho]$ is a singleton), the expression is interpreted as $1 \cdot \prod_{i=1}^{N} \cdots = 1$.

Let $\Psi := \Psi_o \cup \{\emptyset, \iota\}$ be the set of output symbols associated with the events, where $\Psi_o, \emptyset, \iota$ are respectively the observable symbols, unobservable symbol, and starting signal as stated before. Then the sequence of timed output symbols produced by a trajectory is defined as follows:

DEFINITION 5. *Given a trajectory $\rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N}$, the sequence of timed output symbols produced by ρ is*

$$s = (\Delta^0, \psi^0), (\Delta^1, \psi^1) \cdots = \{(\Delta^i, \psi^i)\}_{i=0}^{N},$$

where $(\Delta^0, \psi^0) = (0, \iota)$, and for all $i \geq 1$, $\Delta^i = \tau^{i-1}$, $\psi^i \in \Psi$ is the output symbol associated with $e^i \in E$. For convenience, we define a set of labels $\Sigma := \mathbb{R}_{\geq 0} \times \Psi$, and refer to a sequence of timed output symbols as a label sequence.

2.2 Abstraction

In this section, we construct a system abstraction $\hat{H}$ that helps to analyze the probabilistic fault diagnosability of H and construct a diagnoser. The main result is that $\hat{H}$ can be constructed using finitely many simulated trajectories of H. To that end, we make extensive use of results reported in [11]. The details of [11] are not presented in this paper due to space limitation.

The algorithm in [11] only considers an event with unique target location. It computes robust neighborhoods around the (reset) initial continuous states of any trajectory $\hat{\rho} = \{(\hat{e}^i, \hat{\ell}^i, \hat{x}^i, \hat{\tau}^i)\}_{i=0}^{\hat{N}}$ that has been simulated for finite time horizon, denoted as $Robust(\hat{x}^i)$. The robust neighborhoods computed with a parameter ϵ has the following property [11], where $d_{\mathfrak{R}}$ is a time metric:

- For any (ℓ^0, x^0) that satisfies $\ell^0 = \hat{\ell}^0, x^0 \in Robust(\hat{x}^0)$, there exists a trajectory $\rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N}$ initiated from (ℓ^0, x^0), such that $N = \hat{N}$, $e^i = \hat{e}^i$, $\ell^i = \hat{\ell}^i$, $x^i \in Robust(\hat{x}^i)$, $d_{\mathfrak{R}}(\tau^i, \hat{\tau}^i) \leq \epsilon$ for all $i \in \{0, \ldots, N\}$.

The algorithm can be easily extended to handle an event with multiple target locations: We simulate a trajectory tree $[\hat{\rho}]$ from $(\hat{\ell}^0, \hat{x}^0) \in L^0 \times X^0$ for the time horizon $[0, t_{end}]$, and compute robust neighborhoods for each path of $[\hat{\rho}]$. Then for any (ℓ^0, x^0) such that $\ell^0 = \hat{\ell}^0, x^0 \in \cap_{\hat{\rho} \in [\hat{\rho}]} Robust(\hat{x}^0)$, a set of trajectories $[\rho]$ can be obtained: Any $\rho \in [\rho]$ is initiated from (ℓ^0, x^0) and satisfies the robust neighborhood property above with respect to some $\hat{\rho} \in [\hat{\rho}]$; for each $\hat{\rho} \in [\hat{\rho}]$, only one such ρ needs to be included in $[\rho]$.

By Def. 3 and the robust neighborhood property, clearly $[\rho]$ is a trajectory tree of (ℓ^0, x^0). We thus have the proposition as below:

PROPOSITION 2. *Given a trajectory tree $[\hat{\rho}]$ of $(\hat{\ell}^0, \hat{x}^0)$, for any (ℓ^0, x^0) that satisfies $\ell^0 = \hat{\ell}^0, x^0 \in \cap_{\hat{\rho}\in[\hat{\rho}]} Robust(\hat{x}^0)$, there is a trajectory tree $[\rho]$ of (ℓ^0, x^0), such that the paths of $[\rho]$ and $[\hat{\rho}]$ have a one-to-one correspondence $\mathcal{C} : [\rho] \to [\hat{\rho}]$. For $\rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N} \in [\rho]$, $\hat{\rho} = \{(\hat{e}^i, \hat{\ell}^i, \hat{x}^i, \hat{\tau}^i)\}_{i=0}^{\hat{N}} \in [\hat{\rho}]$, $\hat{\rho} = \mathcal{C}(\rho)$, the following properties hold:*

1. *$N = \hat{N}$, $e^i = \hat{e}^i$, $\ell^i = \hat{\ell}^i$, $x^i \in Robust(\hat{x}^i)$, $d_{\Re}(\tau^i, \hat{\tau}^i) \leq \epsilon$ for all $i \in \{0, \ldots, N\}$.*
2. *Let κ_1 be the Lipschitz constant in Assumption 5. Let $\hat{e}^i = (\hat{\ell}^{i-1}, [\hat{\ell}^i], \hat{g}^i, \hat{r}^i, \hat{p}^i)$ be the i^{th} event triggered by $\hat{\rho}$, $\hat{r}^{i(-1)}(\hat{\ell}^i, Robust(\hat{x}^i))$ be the inverse image of $Robust(\hat{x}^i)$ with respect to the map $\hat{r}^i(\hat{\ell}^i, \cdot) : \hat{g}^i \to Inv(\hat{\ell}^i)$. Given $\kappa_2 > 0$ such that $\|x - \xi_{\hat{\ell}^{i-1}}(\hat{\tau}^{i-1}, \hat{x}^{i-1})\| \leq \kappa_2$ for all $x \in \hat{r}^{i(-1)}(\hat{\ell}^i, Robust(\hat{x}^i))$ and $i \geq 1$, the probability mass of ρ can be estimated by*

$$P(\rho) \in [\prod_{i=1}^{N} \max\{0, \hat{\beta}^i - \kappa\}, \prod_{i=1}^{N} \min\{1, \hat{\beta}^i + \kappa\}], \quad (2)$$

where $\kappa := \kappa_1\kappa_2, \hat{\beta}^i := \hat{p}^i(\hat{\ell}^i, \xi_{\hat{\ell}^{i-1}}(\hat{\tau}^{i-1}, \hat{x}^{i-1}))$. When $N = 0$ ($[\rho]$ is a singleton), the expression is interpreted as $1 \cdot \prod_{i=1}^{N} \cdots = 1$.

PROOF. *Property 1 follows from preceding discussion. To prove Property 2, note that by Property 1, $x^i \in Robust(\hat{x}^i)$, so $\xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1}) \in \hat{r}^{i(-1)}(\hat{\ell}^i, Robust(\hat{x}^i))$. It follows that $\|\xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1}) - \xi_{\hat{\ell}^{i-1}}(\hat{\tau}^{i-1}, \hat{x}^{i-1})\| \leq \kappa_2$. Since $\ell^i = \hat{\ell}^i$, $\|p^i(\ell^i, \xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1})) - \hat{p}^i(\hat{\ell}^i, \xi_{\hat{\ell}^{i-1}}(\hat{\tau}^{i-1}, \hat{x}^{i-1}))\| \leq \kappa_1\kappa_2 = \kappa$. By using Eq. (1), clearly Eq. (2) holds.* □

For convenience we call $\cap_{\hat{\rho}\in[\hat{\rho}]} Robust(\hat{x}^0)$ the robust neighborhood around $\hat{x}^0$ from this point on. As the next step, we need to cover the initial set with the union of robust neighborhoods computed around more simulated initial states [11].

The system abstraction is constructed as $\hat{H} = (L \times X, \hat{L}^0 \times \hat{X}^0, D, E, Inv)$, where $\hat{L}^0 \times \hat{X}^0$ consists of the finitely many simulated initial states whose robust neighborhoods cover $L^0 \times X^0$. Then for any $(\ell^0, x^0) \in L^0 \times X^0$, there is $(\hat{\ell}^0, \hat{x}^0) \in \hat{L}^0 \times \hat{X}^0$, and respective trajectory trees $[\rho], [\hat{\rho}]$, such that the properties in Proposition 2 hold.

Note that given a trajectory tree $[\rho]$ and a subset of paths $\{\rho_k = \{(e_k^i, \ell_k^i, x_k^i, \tau_k^i)\}_{i=0}^{N_k}\}_{k=1}^{K} \subset [\rho], K > 1$, if there is $N^* \geq 0$ such that $(e_k^i, \ell_k^i, x_k^i, \tau_k^i) = (e_{k'}^i, \ell_{k'}^i, x_{k'}^i, \tau_{k'}^i)$ for all $k, k' \in [1, K], i \in [0, N^*]$, and also $\sum_{k=1}^{K} \prod_{i=N^*+1}^{N_k} \beta_k^i = 1$ holds, where $\beta_k^i := p_k^i(\ell_k^i, \xi_{\ell_k^{i-1}}(\tau_k^{i-1}, x_k^{i-1}))$, then clearly the total probability mass of $\{\rho_k\}_{k=1}^{K}$ can be expressed as $\prod_{i=1}^{N^*} \beta^i$ with $\beta^i = \beta_k^i$ for any k. In this case, we say that the trajectories $\{\rho_k\}_{k=1}^{K}$ can be *combined* to the trajectory $\rho^* = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N^*}$, where the sequence $\{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N^*}$ is given by $\{(e_k^i, \ell_k^i, x_k^i, \tau_k^i)\}_{i=0}^{N^*}$ for any k.

As a more convenient way of checking, if $\{\rho_k\}_{k=1}^{K} \subset [\rho]$ consists of all the trajectories $\rho \in [\rho]$ such that ρ^* is a prefix of ρ (see the definition of sub-trajectories following Def. 2), then $\{\rho_k\}_{k=1}^{K}$ can be combined to ρ^*.

COROLLARY 1. *Let $[\rho], [\hat{\rho}]$ be the trajectory trees in Prop. 2, $[\hat{\rho}]' \subset [\hat{\rho}]$ be a subset of paths. Let $\mathcal{C}^{-1}$ be the inverse map of $\mathcal{C}$ in Prop. 2, and $\mathcal{C}^{-1}([\hat{\rho}'])$ be the image of $[\hat{\rho}]'$ for $\mathcal{C}^{-1}$. Let $\mathcal{B}([\hat{\rho}]')$ denote a set of trajectories generated by combining some paths in $[\hat{\rho}]'$. For an arbitrary $\rho^* \in \mathcal{B}([\hat{\rho}]')$, denote ρ^* as $\{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N^*}$, and $\beta^i := p^i(\ell^i, \xi_{\ell^{i-1}}(\tau^{i-1}, x^{i-1}))$ for all $i \in [1, N^*]$, $\kappa := \kappa_1\kappa_2$. Then the total probability mass of $\mathcal{C}^{-1}([\hat{\rho}]') \subset [\rho]$ can be estimated by:*

$$\sum_{\rho\in\mathcal{C}^{-1}([\hat{\rho}]')} P(\rho) \geq \sum_{\rho^*\in\mathcal{B}([\rho]')} \prod_{i=1}^{N^*} \max\{0, \beta^i - \kappa\}, \quad (3)$$

$$\sum_{\rho\in\mathcal{C}^{-1}([\hat{\rho}]')} P(\rho) \leq \sum_{\rho^*\in\mathcal{B}([\rho]')} \prod_{i=1}^{N^*} \min\{1, \beta^i + \kappa\}]. \quad (4)$$

PROOF. *If $\{\hat{\rho}_k\}_{k=1}^{K} \subset [\hat{\rho}]'$ can be combined to $\rho^* \in \mathcal{B}([\hat{\rho}]')$, the total probability mass of $\{\hat{\rho}_k\}_{k=1}^{K} \subset [\hat{\rho}]'$ can be expressed as $\prod_{i=1}^{N^*} \beta^i$. By Prop. 2, clearly $\{\mathcal{C}^{-1}(\hat{\rho}_k)\}_{k=1}^{K}$ can also be combined, and the total probability mass of $\{\mathcal{C}^{-1}(\hat{\rho}_k)\}_{k=1}^{K}$ is within the interval $[\prod_{i=1}^{N^*} \max\{0, \beta^i - \kappa\}, \prod_{i=1}^{N^*} \min\{1, \beta^i + \kappa\}]$. So the result follows.* □

The system abstraction $\hat{H}$ has finitely many initial states. Its diagnosability can be easily analyzed and used to derive the diagnosability of H. We see this in the next section.

3. PROBABILISTIC DIAGNOSABILITY

3.1 Projected Label Sequences

In Section 3, we investigate the problem of diagnosing faults for hybrid systems with probabilistic reset without directly observing the trajectories.

During the system operation, one can only observe a sequence of timed output symbols, which we call a label sequence by Def. 5. Due to the unobservable events, this observed label sequence may be different from the original one produced by the trajectory. Hence, we introduce the definition of *projected label sequences*:

DEFINITION 6. *Let Σ^* denote the set of all the label sequences generated over Σ. Let $s = \{(\Delta^i, \psi^i)\}_{i=0}^{N} \in \Sigma^*$ be a label sequence, and $\Pi : \Sigma^* \to \Sigma^*$ be a single-valued projection map. Then $\pi := \Pi(s)$ is called the projected label sequence of s through the map Π.*

We define the projection map Π that absorbs all the labels with the unobservable output symbol $\emptyset$ into the first observable label that follows, while leaves the rest of the labels unchanged. For instance, $(\Delta^0, \psi^0), (\Delta^1, \emptyset), (\Delta^2, \psi^2)$ is projected to $(\Delta^0, \psi^0), (\Delta^1 + \Delta^2, \psi^2)$. If a trajectory has consecutive unobservable output symbols at its end, then the unobservable end is abandoned in the projected label sequence. Formally the projection map is defined as below:

DEFINITION 7. *Given a label sequence $s = \{(\Delta^i, \psi^i)\}_{i=0}^{N}$, whose observable output symbol sequence is then written as $\{\psi^{i_n}\}_{n=0}^{N'} \subset \{\Psi_o, \iota\}, N' \leq N$, define the projection map:*

$$\Pi(s) = \{(\Delta'^n, \psi'^n)\}_{n=0}^{N'},$$

where $(\Delta'^0, \psi'^0) = (\Delta^0, \psi^0) = (0, \iota)$, *and* $\forall n \in [1, N']$, $(\Delta'^n, \psi'^n) = (\sum_{i=i_{n-1}+1}^{i_n} \Delta^i, \psi^{i_n})$.

Projected label sequences are the only accessible information for system diagnosis. They contain the following aspects of information:

1. Fault diagnosis begins from the starting signal $\psi'^0 = \iota$. So we assume the system operation is clear of faults before the starting signal.
2. Each observable output symbol $\psi'^n, n \geq 1$ should be generated Δ'^n time units later than the preceding one; in the meantime, no observable output symbol can be generated.

In practice, it is often the case that the intervals between the observed output symbols cannot be measured precisely. Given δ_m as the uncertainty parameter, the measurement $\bar{\Delta}'^n$ of Δ'^n must be located in a $\delta_m/2$ neighborhood of Δ'^n.

In what follows, a metric on Σ^* is defined in such a way that the distance only depends on the time sequences if the output symbol sequences of two label sequences are the same, and raised to infinity otherwise. This is motivated by the application of diagnosing faults by observing the projected label sequences. Unlike dwell time, which may be measured with uncertainty, different output symbols are assumed to be readily differentiable from each other.

Definition 8. *Given* $s_1 = \{\Delta_1^i, \psi_1^i\}_{i=0}^{N_1}, s_2 = \{\Delta_2^i, \psi_2^i\}_{i=0}^{N_2}$,

$$d_{\Sigma^*}(s_1, s_2) \triangleq \begin{cases} \sup_i d_{\Re}(\Delta_1^i, \Delta_2^i) & \text{if } N_1 = N_2, \\ & \text{and } \forall i \in [0, N_1], \psi_1^i = \psi_2^i; \\ \infty & \text{otherwise.} \end{cases}$$

Definition 9 (Relative Time Metric).

$$d_{\Re}(t_1, t_2) \triangleq \begin{cases} 0 & \text{if } t_1 = t_2 = 0, \\ \frac{|t_1 - t_2|}{t_1 + t_2} & \text{otherwise.} \end{cases}$$

Proposition 3. $d_{\Re}$ *is a metric on* $\mathbb{R}_{\geq 0}$. *The closed ball* $\{t | d_{\Re}(t_1, t) \leq \epsilon\}$ *is given by the interval* $[t_1 \frac{1-\epsilon}{1+\epsilon}, t_1 \frac{1+\epsilon}{1-\epsilon}]$.

Proof. *Obviously,* $d_{\Re}(t_1, t_2) = d_{\Re}(t_2, t_1)$, $d_{\Re}(t_1, t_2) \geq 0$, *and* $d_{\Re}(t_1, t_2) = 0$ *if and only if* $t_1 = t_2$. *Assume* $t_1 \geq t_2 \geq t_3$. *Then*

$$\begin{aligned} & d_{\Re}(t_1, t_2) + d_{\Re}(t_2, t_3) - d_{\Re}(t_1, t_3) \\ = \; & \frac{t_1 - t_2}{t_1 + t_2} + \frac{t_2 - t_3}{t_2 + t_3} - \frac{t_1 - t_3}{t_1 + t_3} \\ = \; & \frac{(t_1 - t_3)(t_1 - t_2)(t_2 - t_3)}{(t_1 + t_2)(t_2 + t_3)(t_1 + t_3)} \\ \geq \; & 0. \end{aligned}$$

Using similar arguments we can prove $d_{\Re}(t_1, t_2) + d_{\Re}(t_1, t_3) \geq d_{\Re}(t_2, t_3)$ *and* $d_{\Re}(t_1, t_3) + d_{\Re}(t_2, t_3) \geq d_{\Re}(t_1, t_2)$.

To compute the ball centered at t_1 *with radius* ϵ:
If $t \geq t_1$, *then* $d_{\Re}(t_1, t) \leq \epsilon \Rightarrow \frac{t - t_1}{t_1 + t} \leq \epsilon \Rightarrow t \leq t_1 \frac{1+\epsilon}{1-\epsilon}$.
If $t < t_1$, *then* $d_{\Re}(t_1, t) \leq \epsilon \Rightarrow \frac{t_1 - t}{t_1 + t} \leq \epsilon \Rightarrow t \geq t_1 \frac{1-\epsilon}{1+\epsilon}$. □

This metric is relative in the sense that it depends on the elapsed time, i.e., the distance between a 1-second interval and a 1.05-second one is the same as the distance between a 100-second interval and a 105-second one. The motivation of using the relative time metric rather than the more intuitive Euclidean metric is the proof of Prop. 4 (Section 3.3).

3.2 Definition of $(\delta_d, \delta_m, \alpha)$-Diagnosability

Consider a hybrid system with probabilistic reset $H = (L \times X, L^0 \times X^0, D, E, Inv)$, and the projection map Π defined in Section 3.1. Let $E^f \subset E \cup \{e^0\}$ be the set of events that model a fault. We call E^f the faulty events, whose elements can be partitioned into M disjoint subsets $\bigcup_{j=1}^{M} E_j^f = E^f$. Each faulty subset E_j^f corresponds to a type of fault $F_j \in \mathbb{F} := \{F_1, \ldots, F_M\}$. The probabilistic fault diagnosability is proposed in this section, which takes into account the delay in discriminating a fault since it first occurs, as well as the measurement uncertainty of time intervals.

We start from defining the trajectories that triggers some faulty event and then keeps flowing for enough long time:

Definition 10. *A trajectory* $\rho = \{(e^i, \ell^i, x^i, \tau^i)\}_{i=0}^{N}$ *is* $F_j\delta$*-faulty if and only if there exists a finite index* $i^f \in [0, N]$ *such that:*

1. $\forall i < i^f$, $e^i \notin E_j^f$;
2. $e^{i^f} \in E_j^f$;
3. $\sum_{i=i^f}^{N} \tau^i \geq \delta$.

In the special case where $\delta = 0$, *i.e., only the occurrence of* F_j *is required, the term* F_j*-faulty is used. If a trajectory is free of any fault, then it is called normal.*

Later, an approach is presented that analyzes the finite-horizon fault diagnosability of H through the analysis of the abstraction $\hat{H}$ (Section 2.2). Recall that the initial set of $\hat{H}$ is a finite set; in the construction of $\hat{L}^0 \times \hat{X}^0$, a trajectory tree $[\hat{\rho}]$ has been simulated for the horizon $[0, t_{end}]$ for each $(\hat{\ell}^0, \hat{x}^0) \in \hat{L}^0 \times \hat{X}^0$. We use $\mathcal{J}(H)$ to denote the set of trajectories that are of interest in the finite-horizon fault diagnosis. $\mathcal{J}(H)$ consists of all the trajectory trees $[\rho]$ obtained from $[\hat{\rho}]$ (see Proposition 2) as (ℓ^0, x^0) varies in $L^0 \times X^0$, and the sub-trajectories of these trajectory trees. Accordingly, the set of trajectories $\mathcal{J}(\hat{H})$ that are of interest for $\hat{H}$ is specified as the collection of all the simulated trajectory trees $[\hat{\rho}]$ and their sub-trajectories.

To avoid ambiguity, from this point on $[\rho] \in \mathcal{J}(H), [\hat{\rho}] \in \mathcal{J}(\hat{H})$ only refer to trajectory trees that extend to the end of time horizon of interest, rather than any trajectory tree formed by their sub-trajectories.

Based on the discriminability of $F_j\delta$-faulty trajectories from normal trajectories and faulty trajectories of other faults, the probabilistic fault detectability and isolability can be defined as below:

DEFINITION 11. *Given a hybrid automaton H and a faulty event set $E^f = \{E_1^f, \ldots, E_M^f\}$, let $\mathcal{F}_j(H) \subset \mathcal{J}(H)$ denote the set of all the F_j-faulty trajectories, and $\mathcal{F}_j^{\delta}(H) \subset \mathcal{F}_j(H)$ denote the set of all the $F_j\delta$-faulty trajectories. Let s_1, s_2 denote the label sequences produced by the trajectories ρ_1, ρ_2 respectively, and $\pi_1 = \Pi(s_1), \pi_2 = \Pi(s_2)$.*

1. *$\mathbb{F}$ is called $(\delta_d, \delta_m, \alpha)$-detectable if and only if for any trajectory tree $[\rho]$ in $\mathcal{J}(H)$, $\sum_{\rho \in \mathcal{A}_1([\rho])} P(\rho) \leq 1 - \alpha$ holds, where $\mathcal{A}_1([\rho]) \subset [\rho]$ is the set of all paths $\rho \in [\rho]$ such that:*
   - *ρ has a sub-trajectory ρ_1, or $\rho = \rho_1$, which is $F_j\delta_d$-faulty for some fault $F_j \in \mathbb{F}$, while cannot be discriminated from some normal trajectory $\rho_2 \in \mathcal{J}(H)$; formally, $\rho_1 \in \mathcal{F}_j^{\delta_d}(H)$, ρ_2 is normal, and $d_{\Sigma^*}(\pi_1, \pi_2) \leq \delta_m$.*
2. *$\mathbb{F}$ is called $(\delta_d, \delta_m, \alpha)$-isolable if and only if for any trajectory tree $[\rho]$ in $\mathcal{J}(H)$, $\sum_{\rho \in \mathcal{A}_2([\rho])} P(\rho) \leq 1 - \alpha$ holds, where $\mathcal{A}_2([\rho]) \subset [\rho]$ is the set of all paths $\rho \in [\rho]$ such that:*
   - *ρ has a sub-trajectory ρ_1, or $\rho = \rho_1$, which is $F_j\delta_d$-faulty for some fault $F_j \in \mathbb{F}$, while cannot be discriminated from some F_i-faulty trajectory $\rho_2 \in \mathcal{F}_i(H)$, where $F_i \in \mathbb{F}, F_i \neq F_j$; formally, $\rho_1 \in \mathcal{F}_j^{\delta_d}(H)$, $\rho_2 \in \mathcal{F}_i(H)$, $d_{\Sigma^*}(\pi_1, \pi_2) \leq \delta_m$.*

Similarly, the $(\delta_d, \delta_m, \alpha)$-diagnosability that requires both probabilistic fault detectability and isolability is defined:

DEFINITION 12. *Given a hybrid automaton H and a faulty event set $E^f = \{E_1^f, \ldots, E_M^f\}$, the system H is said to be $(\delta_d, \delta_m, \alpha)$-diagnosable if and only if for any trajectory tree $[\rho]$ in $\mathcal{J}(H)$, $\sum_{\rho \in \mathcal{A}([\rho])} P(\rho) \leq 1 - \alpha$ holds, where $\mathcal{A}([\rho]) := \mathcal{A}_1([\rho]) \cup \mathcal{A}_2([\rho])$, $\mathcal{A}_1, \mathcal{A}_2$ are defined in Def. 11.*

The probabilistic diagnosability of faults can be extended from single-fault to multiple-fault cases. For example, given $\mathbb{F} = \{F_1, F_2, F_3\}$ and a combination of faults $f = \{F_1, F_3\} \in 2^{\mathbb{F}}$, define $f\delta$-faulty trajectories as both $F_1\delta$-faulty and $F_3\delta$-faulty, i.e., the intersection of the sets of $F_1\delta$-faulty and $F_3\delta$-faulty trajectories forms the set of $f\delta$-faulty trajectories. Then multiple-fault probabilistic diagnosability can be defined and analyzed based on the discriminability of $f\delta$-faulty trajectories.

3.3 Bridging The Abstraction to Probabilistic Fault Diagnosability Analysis

In this section, we establish a quantitative relation between the probabilistic fault diagnosability of a hybrid system and the abstraction.

PROPOSITION 4. *Given $s_1, s_2 \in \Sigma^*$ and $\pi_1 = \Pi(s_1), \pi_2 = \Pi(s_2)$, we have $d_{\Sigma^*}(\pi_1, \pi_2) \leq d_{\Sigma^*}(s_1, s_2)$.*

PROOF. If $d_{\Sigma^*}(s_1, s_2)$ is finite, then clearly s_1, s_2 have the same output symbol sequence. According to the definition of Π, π_1, π_2 must also have the same output symbol sequence. Moreover, for any $\Delta_1^1, \Delta_2^1, \Delta_1^2, \Delta_2^2, > 0$, $\frac{|\Delta_1^1 + \Delta_1^2 - \Delta_2^1 - \Delta_2^2|}{\Delta_1^1 + \Delta_1^2 + \Delta_2^1 + \Delta_2^2} \leq \frac{|\Delta_1^1 - \Delta_2^1| + |\Delta_1^2 - \Delta_2^2|}{\Delta_1^1 + \Delta_1^2 + \Delta_2^1 + \Delta_2^2} \leq \max\{\frac{|\Delta_1^1 - \Delta_2^1|}{\Delta_1^1 + \Delta_2^1}, \frac{|\Delta_1^2 - \Delta_2^2|}{\Delta_1^2 + \Delta_2^2}\}$. Therefore, by combining some labels' dwell time distance after the projection, the supremum of dwell time distance over all the labels remains unchanged or becomes smaller.

Thus, we obtain $d_{\Sigma^*}(\pi_1, \pi_2) \leq d_{\Sigma^*}(s_1, s_2)$. □

THEOREM 1. *Given $\hat{\delta}_d, \hat{\delta}_m$, the maximum value of $\hat{\alpha}$ such that $\hat{H}$ is $(\hat{\delta}_d, \hat{\delta}_m, \hat{\alpha})$-diagnosable can be calculated by using the formula:*

$$\begin{aligned} \hat{\alpha} &= \min_{[\hat{\rho}] \subset \mathcal{J}(\hat{H})} \sum_{\hat{\rho} \in [\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}])} P(\hat{\rho}) \\ &= \min_{[\hat{\rho}] \subset \mathcal{J}(\hat{H})} \sum_{\hat{\rho} \in [\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}])} \prod_{i=1}^{\hat{N}} \hat{\beta}^i; \\ &= \min_{[\hat{\rho}] \subset \mathcal{J}(\hat{H})} \sum_{\rho^* \in \mathcal{B}([\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}]))} \prod_{i=1}^{N^*} \beta^i; \end{aligned} \tag{5}$$

and H is $(\delta_d^*, \delta_m^*, \alpha^*)$-diagnosable with

$$\delta_d^* = \hat{\delta}_d \frac{1+\epsilon}{1-\epsilon}, \tag{6}$$

$$\delta_m^* = \hat{\delta}_m - 2\epsilon, \tag{7}$$

$$\alpha^* = \min_{[\hat{\rho}] \subset \mathcal{J}(\hat{H})} \sum_{\rho^* \in \mathcal{B}([\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}]))} \prod_{i=1}^{N^*} \max\{0, \beta^i - \kappa\}, \tag{8}$$

where $\mathcal{A}([\hat{\rho}]) := \mathcal{A}_1([\hat{\rho}]) \cup \mathcal{A}_2([\hat{\rho}])$ is defined in Def. 11, 12, and $\kappa, \hat{\beta}^i, \mathcal{B}, \beta^i$ are defined in Proposition 2, Corollary 1.

PROOF. *The abstraction $\hat{H}$ constructed in Section 2.2 has finitely many trajectory trees. The calculation of $\hat{\alpha}$ directly follows from Def. 11 and 12.*

Since $L^0 \times X^0$ is covered by the robust neighborhoods of $\hat{L}^0 \times \hat{X}^0$, for any trajectory tree $[\rho]$ in $\mathcal{J}(H)$, there exists a trajectory tree $[\hat{\rho}]$ in $\mathcal{J}(\hat{H})$, such that the paths of $[\rho]$ and $[\hat{\rho}]$ have a one-to-one correspondence $\mathcal{C} : [\rho] \to [\hat{\rho}]$ with the properties in Prop. 2. Let ρ be an arbitrary path of $[\rho]$, and $\hat{\rho} \in [\hat{\rho}]$ be $\mathcal{C}(\rho)$.

Given the parameters $\delta_d = \delta_d^*, \delta_m = \delta_m^*$, suppose $\rho \in \mathcal{A}([\rho])$. This implies, by Def. 11 and 12, the existence of ρ_1 as a sub-trajectory of ρ or $\rho_1 = \rho$, a trajectory $\rho_2 \in \mathcal{J}(H)$, and a fault $F_j \in \mathbb{F}$, such that $\rho_1 \in \mathcal{F}_j^{\delta_d^*}(H)$, ρ_2 is either normal or contained in $\mathcal{F}_i(H)$ for some fault $F_i \neq F_j$, and their projected label sequences satisfy $d_{\Sigma^*}(\pi_1, \pi_2) \leq \delta_m^*$. Denote ρ_1, ρ_2 as

$$\begin{aligned} \rho_1 &= (e_1^0, \ell_1^0, x_1^0, \tau_1^0), \ldots, (e_1^{i^f}, \ell_1^{i^f}, x_1^{i^f}, \tau_1^{i^f}), \\ &\quad \ldots, (e_1^{N_1}, \ell_1^{N_1}, x_1^{N_1}, \tau_1^{N_1}), \\ \rho_2 &= (e_2^0, \ell_2^0, x_2^0, \tau_2^0), \ldots, (e_2^{N_2}, \ell_2^{N_2}, x_2^{N_2}, \tau_2^{N_2}). \end{aligned}$$

Since $\rho_1 \in \mathcal{F}_j^{\delta_d^*}(H)$, we have $\sum_{i=i^f}^{N_1} \tau_1^i \geq \delta_d^*$.

By Prop. 2, there exist $\hat{\rho}_1$ as a sub-trajectory of $\hat{\rho}$ or $\hat{\rho}_1 = \hat{\rho}$, and a trajectory $\hat{\rho}_2 \in \mathcal{J}(\hat{H})$, such that $N_1 = \hat{N}_1$, $e_1^i = \hat{e}_1^i$,

$\ell_1^i = \hat{\ell}_1^i$, $d_{\Re}(\tau_1^i, \hat{\tau}_1^i) \leq \epsilon$ *for all* $i \in [0, N_1]$; $N_2 = \hat{N}_2$, $e_2^i = \hat{e}_2^i$, $\ell_2^i = \hat{\ell}_2^i$, $d_{\Re}(\tau_2^i, \hat{\tau}_2^i) \leq \epsilon$ *for all* $i \in [0, N_2]$. *Denote* $\hat{\rho}_1, \hat{\rho}_2 \in \mathcal{J}(\hat{H})$ *as*

$$\begin{aligned}\hat{\rho}_1 &= (\hat{e}_1^0, \hat{\ell}_1^0, \hat{x}_1^0, \hat{\tau}_1^0), \ldots, (\hat{e}_1^{i^f}, \hat{\ell}_1^{i^f}, \hat{x}_1^{i^f}, \hat{\tau}_1^{i^f}),\\ &\quad \ldots, (\hat{e}_1^{\hat{N}_1}, \hat{\ell}_1^{\hat{N}_1}, \hat{x}_1^{\hat{N}_1}, \hat{\tau}_1^{\hat{N}_1}),\\ \hat{\rho}_2 &= (\hat{e}_2^0, \hat{\ell}_2^0, \hat{x}_2^0, \hat{\tau}_2^0), \ldots, (\hat{e}_2^{\hat{N}_2}, \hat{\ell}_2^{\hat{N}_2}, \hat{x}_2^{\hat{N}_2}, \hat{\tau}_2^{\hat{N}_2}).\end{aligned}$$

Clearly, ρ_1 *and* $\hat{\rho}_1$ *have the same event sequence,* ρ_2 *and* $\hat{\rho}_2$ *have the same event sequence. It follows that* $\hat{\rho}_1 \in \mathcal{F}_j(\hat{H})$, $\hat{\rho}_2$ *is either normal or contained in* $\mathcal{F}_i(\hat{H})$. *By Def. 8, we also have*

$$d_{\Sigma^*}(s_1, \hat{s}_1) \leq \epsilon, d_{\Sigma^*}(s_2, \hat{s}_2) \leq \epsilon. \tag{9}$$

Let $\pi_1, \pi_2, \hat{\pi}_1, \hat{\pi}_2$ *be the projected label sequences produced respectively by* $\rho_1, \rho_2, \hat{\rho}_1, \hat{\rho}_2$. *By the triangle inequality, Prop. 4 and Eq.* (9)*, the following holds:*

$$\begin{aligned}d_{\Sigma^*}(\hat{\pi}_1, \hat{\pi}_2) &\leq d_{\Sigma^*}(\pi_1, \hat{\pi}_1) + d_{\Sigma^*}(\pi_1, \pi_2) + d_{\Sigma^*}(\pi_2, \hat{\pi}_2)\\ &\leq \delta_m^* + 2\epsilon\\ &= \hat{\delta}_m.\end{aligned}$$

For all $i \in [0, N_1]$, $d_{\Re}(\tau_1^i, \hat{\tau}_1^i) \leq \epsilon$, *namely,* $\tau_1^i \frac{1-\epsilon}{1+\epsilon} \leq \hat{\tau}_1^i \leq \tau_1^i \frac{1+\epsilon}{1-\epsilon}$. *It follows that*

$$\sum_{i=i^f}^{\hat{N}_1} \hat{\tau}_1^i \geq \sum_{i=i^f}^{N_1} \tau_1^i \frac{1-\epsilon}{1+\epsilon} \geq \delta_d^* \frac{1-\epsilon}{1+\epsilon}. \tag{10}$$

By definition, $\hat{\rho}_1 \in \mathcal{F}_j^{\delta_d}(\hat{H})$ *with* $\delta_d = \delta_d^* \frac{1-\epsilon}{1+\epsilon} = \hat{\delta}_d$, *and thus* $\hat{\rho} \in \mathcal{A}([\hat{\rho}])$ *for the parameters* $\delta_d = \hat{\delta}_d, \delta_m = \hat{\delta}_m$.

Therefore, $\hat{\rho} \in [\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}])$ *implies* $\rho = \mathcal{C}^{-1}(\hat{\rho}) \in [\rho] \setminus \mathcal{A}([\rho])$, *which means* $\mathcal{C}^{-1}([\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}])) \subset [\rho] \setminus \mathcal{A}([\rho])$. *It follows from Corollary 1 that*

$$\sum_{\rho \in [\rho] \setminus \mathcal{A}([\rho])} P(\rho) \geq \sum_{\rho^* \in \mathcal{B}([\hat{\rho}] \setminus \mathcal{A}([\hat{\rho}]))} \prod_{i=1}^{N^*} \max\{0, \beta^i - \kappa\}. \tag{11}$$

□

3.4 Diagnosers

The abstraction $\hat{H}$ constructed in Section 2.2 has finitely many trajectories extending to the time horizon of interest. In this section, a diagnoser based on $\hat{H}$ is constructed, which works in the following way: It stores a finite list of candidate paths whose event sequences can be triggered by H, and keeps narrowing down the list by observing the timed output symbol sequence of H until a decision is made. When a decision has to be made without certainty, the diagnoser can make choices from the candidates based on their probability conditioning on the current observation [7].

Suppose the delay parameter and measurement uncertainty for H are given by δ_d^*, δ_m^*. That is, the measurements of time intervals must lie inside a ball of radius $\frac{1}{2}\delta_m^*$ centered at the true value. We construct $\hat{H}$ with a parameter ϵ, and compute $\hat{\delta}_d = \delta_d^* \frac{1-\epsilon}{1+\epsilon}, \hat{\delta}_m = \delta_m^* + 2\epsilon$, and $\hat{\alpha}$ as Eq. (5). By Theorem 1, $\hat{H}$ is $(\hat{\delta}_d, \hat{\delta}_m, \hat{\alpha})$-diagnosable, H is $(\delta_d^*, \delta_m^*, \alpha^*)$-diagnosable with α^* computed from Eq. (8).

The diagnoser is then constructed as a hybrid automaton $H_d = (L_d \times X_d, L_d^0 \times X_d^0, D_d, E_d, Inv_d)$, where the sematics is the same as Def. 1 except that each event has a unique target location.

- Let $\{\hat{\rho}_k\}_{k=1}^K$ be the collection of all the paths of all $[\hat{\rho}] \in \mathcal{J}(\hat{H})$, where $\hat{\rho}_k = \{(\hat{e}_k^i, \hat{\ell}_k^i, \hat{x}_k^i, \hat{\tau}_k^i)\}_{i=0}^{\hat{N}_k}$. The state space of the diagnoser is defined as $L_d := 2^{\{1,\ldots,K\}} \times \{0, 1, 2, \ldots\}$, $X_d := \mathbb{R}$. Clearly, each location $\ell_d \in L_d$ is a set of pairs $(k, n), k \in \{1, \ldots, K\}, n \in \{0, 1, 2, \ldots\}$.
- Let $\hat{s}_k = \{(\hat{\Delta}_k^i, \hat{\psi}_k^i)\}_{i=0}^{\hat{N}_k}$ be produced by $\hat{\rho}_k$, and $\hat{\psi}_k^{i_{(k,n)}}$ be the n^{th} observable output symbol of $\hat{\rho}_k$ following the starting signal ($i_{(k,n)}$ is the index of $\hat{\psi}_k^{i_{(k,n)}}$ in $\hat{s}_k$). We define the set of fault labels $W := 2^{\mathbb{F}}$, where $\mathbb{F} := \{F_1, \ldots, F_M\}$ are the modeled M types of faults. Each pair $(k, n), n \in \{0, 1, \ldots, \hat{N}_k\}$ possesses a complete fault label $w_{(k,n)}^{all} \in W$ as the collection of all faults made by $\hat{\rho}_k$, a current fault label $w_{(k,n)}^{now} \subset w_{(k,n)}^{all}$ as the set of faults made by event sequence of $\hat{\rho}_k$ upto index $i_{(k,n)}$, and a overdue fault label $w_{(k,n)}^{due} \subset w_{(k,n)}^{now}$ such that $\sum_{i=i^f+1}^{i_{(k,n)}} \hat{\Delta}_k^i \geq \hat{\delta}_d$.
- Given a location ℓ_d, define the conditional probability scores $y_{(k,n)}$ for each pair $(k, n) \in \ell_d$:

$$y_{(k,n)} := \frac{P(\hat{\rho}_k)}{\Sigma_{k' \in [k]'} P(\hat{\rho}_{k'})}, \tag{12}$$

  where $[k]'$ is the subset of $\{1, \ldots, K\}$ such that $k' \in [k]'$ satisfies $(k', n) \in \ell_d$, and $\hat{\rho}_{k'}$ has the same initial state as $\hat{\rho}_k$, i.e., $\hat{\rho}_{k'}, \hat{\rho}_k$ belong to the same trajectory tree.
- $L_d^0 \times X_d^0 := (\{1, \ldots, K\}, 0) \times \{0\}$.
- $D_d : \dot{x} = 1$ for all the locations $\ell_d \in L_d$.
- Let $\hat{\pi}_k = \{(\hat{\Delta}_k'^n, \hat{\psi}_k^{i_{(k,n)}})\}_{i=0}^{\hat{N}_k'} = \Pi(\hat{s}_k)$, where $\hat{\Delta}_k'^n = \sum_{i=i_{(k,n-1)}+1}^{i_{(k,n)}} \hat{\Delta}_k^i$. Whenever a timed output symbol $(\bar{\Delta}, \psi)$ is observed from H, the diagnoser H_d will reset its location from ℓ_d to ℓ_d' by triggering an event. The target ℓ_d' is defined as the set of all pairs (k, n) such that $(k, n-1) \in \ell_d$, $\psi = \hat{\psi}_k^{i_{(k,n)}}$, and $\bar{\Delta} \in \mathbb{B}(\hat{\Delta}_k'^n, \frac{1}{2}\hat{\delta}_m)$, where $\mathbb{B}(\hat{\Delta}_k'^n, \frac{1}{2}\hat{\delta}_m)$ is the relative time metric ball centered at $\hat{\Delta}_k'^n$ with radius $\frac{1}{2}\hat{\delta}_m$. The event also resets the continuous state to 0.
- $Inv_d(\ell_d) := \mathbb{R}$ for all $\ell_d \in L_d$.

The diagnoser H_d operates along with H. Clearly, H_d is deterministic. The continuous state x_d is the time. The discrete state ℓ_d updates according to the observed events of H with their timing, where measurement uncertainty exists. Thus, despite the difference between $\bar{\Delta}$ and $\hat{\Delta}_k'^n$, we still consider $\hat{\rho}_k$ as a candidate path and include (k, n) in the updated diagnoser state, as long as the difference in time intervals is bounded by $\frac{1}{2}\hat{\delta}_m = \frac{1}{2}\delta_m^* + \epsilon$, and the output symbol matches the observation completely.

Whenever H_d reaches a location ℓ_d such that $w_{(k,n)}^{now} \subset \mathbb{F}$ for all $(k, n) \in \ell_d$, a fault detection alarm should be activated

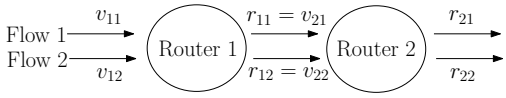

Figure 1: A network model with two routers and two flows. Each router has one buffer. The package loss event is observable at only one of the routers.

with certainty. In particular, if $w^{now}_{(k,n)} \ni F_j$ (or $w^{now}_{(k,n)} = \{F_j\}$) for all $(k,n) \in \ell_d$, the fault F_j is detected (or isolated) with certainty.

The diagnoser could also encounter an undesired case where for $(k_1,n),(k_2,n) \in \ell_d$, the overdue fault label $w^{due}_{(k_1,n)}$ contains F_j, while $w^{now}_{(k_2,n)} \neq \{F_j\}$, meaning that F_j cannot be detected or isolated in time. According to Def. 11, 12, $(1-\alpha^*)$ overestimates the probability that some fault will be incurred for at least δ^*_d time units while fail to be detected or isolated. Thus, if we stick to the rule that a detection/isolation alarm is raised only when a fault can be diagnosed with certainty, the overall detection failure (false negative) and isolation failure rate should not exceed the threshold $(1-\alpha^*)$.

The diagnoser may run into the undesired case mentioned above ($(k_1,n),(k_2,n) \in \ell_d, w^{due}_{(k_1,n)} \ni F_j, w^{now}_{(k_2,n)} \neq \{F_j\}$) with probability $1-\bar{a}^* > 1-a^*$, since the probability mass of $\hat{\rho}_{k_2}$ is not necessarily counted in the estimation of $(1-\alpha^*)$. We can easily modify Def. 11, 12 to characterize $(1-\bar{\alpha}^*)$ instead, and all the results including Theorem 1 still hold. The proof is basically the same and not discussed.

If the diagnoser has to make diagnosis in an ambiguous state, the conditional probability scores $y_{(k,n)}$ can be used, which represent the probability of $\hat{\rho}_k$ conditioning on the observed timed events. For example, given $\ell_d = \{(k_1,n),(k_2,n)\}$, $w^{now}_{(k_1,n)} = \{F_j\}, w^{now}_{(k_2,n)} = \{\}$, we can adopt the strategy that F_j alarm is raised if $y_{(k_1,n)} > y_{(k_2,n)}$. The conditional probability here is defined with respect to a trajectory tree $[\hat{\rho}]$ as before, and we may have multiple $[\hat{\rho}]$ involved in ℓ_d. Thus, the scores $y_{(k_1,n)} > y_{(k_2,n)}$ does not necessarily mean that $\hat{\rho}_{k_1}$ is more likely than $\hat{\rho}_{k_2}$. If $\hat{\rho}_{k_1}, \hat{\rho}_{k_2}$ live in the same probability space, their conditional probability scores can offer good reference for the diagnoser to make smarter decisions.

REMARK 1. *By the way we define $w^{all}_{(k,n)}$, if the detection (or isolation) alarm is raised as soon as $w^{all}_{(k,n)} \ni F_j$ (or $w^{all}_{(k,n)} = \{F_j\}$) for all $(k,n) \in \ell_d$, then it is possible to prognose F_j before its occurrence.*

3.5 Implementation Example

In this section we discuss the application of our framework in network package loss diagnosis. For a group of networked routers, it is assumed the package loss event is observable at only some of the routers. The problem is to diagnose the congestion of the whole network. In the present work we only consider a simplified model as follows:

Let v_{ij} and r_{ij} denote respectively the arrival rate and transmission rate of Router i, Flow j. See Figure 1. Let BW_i, q_i, Q_i denote respectively the bandwidth, queue length, and buffer size (maximum queue length) of Router i.

The arrival rate v_{1j} to Router 1 equals the sending rate of Flow j from the sender. For simplicity, we assume the flow sending rates increase linearly. When the the total transmission rate of flows is less than the bandwidth, the transmission rate of each flow equals the arrival rate, and the queue remains empty. If the total transmission rate reaches the bandwidth, and is less than the total arrival rate, then the transmission rate of each flow keeps constant, and the queue starts filling. Once the queue is filled, the buffer drops a package and reduces the queue length by the package length M. The dropped package may come from each of the flows, with the probability equal to the fraction of the flow's arrival rate of the total arrival rate. In the simplified model, it is assumed the sender immediately halves the sending rate of a flow when a package loss occurs. The state space equations are listed below.

1. Both q_1, q_2 are empty.

$$\begin{cases} \dot{v}_{11} = K_1, \\ \dot{v}_{12} = K_2, \\ \dot{r}_{11} = K_1, \\ \dot{r}_{12} = K_2, \\ \dot{r}_{21} = K_1, \\ \dot{r}_{22} = K_2, \\ \dot{q}_1 = 0, \\ \dot{q}_2 = 0. \end{cases}$$

2. q_1 is increasing/decreasing, q_2 is empty.

$$\begin{cases} \dot{v}_{11} = K_1, \\ \dot{v}_{12} = K_2, \\ \dot{r}_{11} = 0, \\ \dot{r}_{12} = 0, \\ \dot{r}_{21} = 0, \\ \dot{r}_{22} = 0, \\ \dot{q}_1 = v_{11} + v_{12} - r_{11} - r_{12}, \\ \dot{q}_2 = 0. \end{cases}$$

3. q_1 is empty, q_2 is increasing/decreasing.

$$\begin{cases} \dot{v}_{11} = K_1, \\ \dot{v}_{12} = K_2, \\ \dot{r}_{11} = K_1, \\ \dot{r}_{12} = K_2, \\ \dot{r}_{21} = 0, \\ \dot{r}_{22} = 0, \\ \dot{q}_1 = 0, \\ \dot{q}_2 = r_{11} + r_{12} - r_{21} - r_{22}. \end{cases}$$

4. Both q_1, q_2 are increasing/decreasing.

$$\begin{cases} \dot{v}_{11} = K_1, \\ \dot{v}_{12} = K_2, \\ \dot{r}_{11} = 0, \\ \dot{r}_{12} = 0, \\ \dot{r}_{21} = 0, \\ \dot{r}_{22} = 0, \\ \dot{q}_1 = v_{11} + v_{12} - r_{11} - r_{12}, \\ \dot{q}_2 = r_{11} + r_{12} - r_{21} - r_{22}. \end{cases}$$

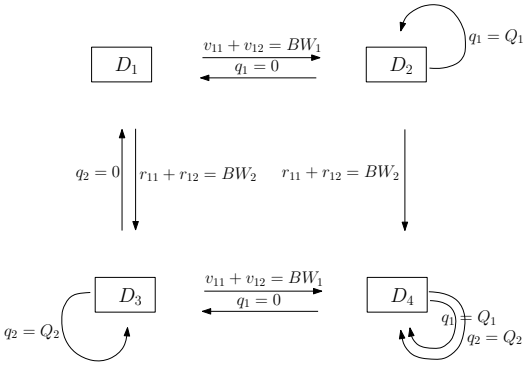

Figure 2: The guards of events are given by the equations. When $v_{i1}+v_{i2}=BW_i$, q_i starts filling; when $q_i=0$, the buffer is emptied; when $q_i=Q_i$, a package is dropped.

The four modes of dynamics listed above are referred to as D_1, D_2, D_3, D_4 respectively. Based on the dynamics, we model the locations and events. As shown in Figure 2, there are three types of events: *bandwidth reached*, *buffer emptied*, *package dropped*.

In particular, when a package loss event occurs, the reset is probabilistic. Either Flow 1 or Flow 2 loses a package and reduces the sending rate to half of the current value. For convenience, besides the locations $\ell_1, \ell_2, \ell_3, \ell_4$ with dynamics D_1 to D_4, we also define locations $\ell_2', \ell_2'', \ell_3', \ell_3'', \ell_4', \ell_4''$. The continuous dynamics of ℓ_i', ℓ_i'' is the same as ℓ_i.

When a package loss event occurs in ℓ_2, the target location can be ℓ_2' or ℓ_2'', respectively meaning Flow 1 or Flow 2 suffers a package loss. The associated reset map $r(\ell_2', x)$ changes v_{11} to $\frac{1}{2}v_{11}$, and $r(\ell_2'', x)$ changes v_{12} to $\frac{1}{2}v_{12}$, where $x = [v_{11}, v_{12}, r_{11}, r_{12}, r_{21}, r_{22}, q_1, q_2]^T$ is the continuous state that triggers the event. The corresponding probability is given by $p(\ell_2', x) = \frac{v_{11}}{v_{11}+v_{12}}$, $p(\ell_2'', x) = \frac{v_{12}}{v_{11}+v_{12}}$. The package loss events in ℓ_3, ℓ_4 are defined similarly.

The probability function p above is Lipschitz continuous with respect to x. We have the Lipschitz constant $\kappa_1 = 1$.

Let $K_1 = 3, K_2 = 1, BW_1 = 8, BW_2 = 4, Q_1 = 5, Q_2 = 7, M = 1$, $L^0 = \{\ell_1\}$, X^0 be the set $\|v_{11}^0 - 1\| \le 0.0001, \|v_{12}^0 - 0.5\| \le 0.0001, r_{i1}^0 = v_{11}^0, r_{i2}^0 = v_{12}^0, q_i^0 = 0$ for any i, and t_{end} be 3.3 time units. By using the Matlab toolbox STRONG [4], we can simulate a trajectory tree $[\hat{\rho}]$ from the initial state $(\hat{\ell}^0, \hat{x}^0) = (\ell_1, [1, 0.5, 1, 0.5, 1, 0.5, 0, 0]^T)$. We enumerate the location sequences, event sequences and label sequences for all the paths of $[\hat{\rho}]$:

path 1: $\ell_1, \ell_3, \ell_4, \ell_4', \ell_4'$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 1), q_2 drops (Flow 1);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.2778, \psi_3)$;

path 2: $\ell_1, \ell_3, \ell_4, \ell_4', \ell_4''$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 1), q_2 drops (Flow 2);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.2778, \psi_3)$;

path 3: $\ell_1, \ell_3, \ell_4, \ell_4'', \ell_4', \ell_4'$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 2), q_1 drops (Flow 1),
q_2 drops (Flow 1);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.1574, \psi_4), (0.1204, \psi_3)$;

path 4: $\ell_1, \ell_3, \ell_4, \ell_4'', \ell_4', \ell_4''$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 2), q_1 drops (Flow 1),
q_2 drops (Flow 2);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.1574, \psi_4), (0.1204, \psi_3)$;

path 5: $\ell_1, \ell_3, \ell_4, \ell_4'', \ell_4'', \ell_4'$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 2), q_1 drops (Flow 2),
q_2 drops (Flow 1);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.1574, \psi_4), (0.1204, \psi_3)$;

path 6: $\ell_1, \ell_3, \ell_4, \ell_4'', \ell_4'', \ell_4''$;
start, BW_2 reached, BW_1 reached,
q_2 drops (Flow 2), q_1 drops (Flow 2),
q_2 drops (Flow 2);
$(0, \iota), (0.625, \psi_1), (1, \psi_2)$,
$(1.2778, \psi_3), (0.1574, \psi_4), (0.1204, \psi_3)$;

where $\psi_1, \psi_2, \psi_3, \psi_4$ are the output symbols associated with the events BW_2 *full*, BW_1 *full*, q_2 *drops*, q_1 *drops*. Suppose the package loss events are observable only for Router 2. Then ι, ψ_3 are the observable output symbols, while the rest are unobservable.

The event q_1 *drops* is a faulty event we want to diagnose. Given any $\delta_d = \hat{\delta}_d \ge 0$, $\delta_m = \hat{\delta}_m \ge 0$, clearly the $F\delta_d$-faulty trajectories Path 3 to 6 cannot be discriminated from Path 1 or 2, since their projected label sequences are exactly the same. Path 3 to 6 have the total probability mass 0.2656, which equals $(1 - \hat{\alpha})$.

The toolbox computes a robust neighborhood around $(\hat{\ell}^0, x^0)$. With the parameter $\epsilon = 0.02$ in relative time metric, κ_2 (see Prop. 2) is calculated to be 0.001. Therefore, as the

initial state varies in the robust neighborhood with non-determinism, the system is $(1.04\hat{\delta}_d, \hat{\delta}_m - 0.04, \hat{\alpha} - 0.001)$-diagnosable (see Theorem 1; for this particular example, $\hat{\delta}_d, \hat{\delta}_m, \delta_d^*, \delta_m^*$ can actually be any non-negative numbers because the abstraction has no faulty trajectory that can be discriminated from normal trajectories). According to the discussion on the diagnoser in Section 3.4, if we raise an alarm only when the running trajectory is completely determined as faulty, 0.2666 is then an overestimation of the detection failure (false negative) rate of the diagnoser.

From the discussion above, the observation of package loss at Router 2 does not provide detectability of package loss at Router 1. In order to diagnose the network congestion, we can observe the package loss at Router 1 instead: Let ι, ψ_4 be observable, and the rest of output symbols be unobservable. We want to diagnose the combination of the faulty events q_1 *drops,* q_2 *drops* (see the discussion of multiple-fault in Section 3.2). Clearly such a combination is diagnosable by observing the timed output symbols. In specific, as long as ψ_4 is observed (q_1 *drops*), then all the routers in the network drop packages.

4. CONCLUSION

We presented an approach to finite-horizon diagnosability analysis for hybrid systems with probabilistic reset. In our notion of probabilistic diagnosability, given the maximum delay for fault diagnosis and the measurement uncertainty of time intervals, the worst-case probability of detecting and isolating the faults is characterized. So it generalizes some existing diagnosability notions [5, 16]. To practically analyze such diagnosability for a hybrid system, we constructed a system abstraction, whose probabilistic diagnosability is proved to be quantitatively related to that of the original system. The abstraction has finitely many initial states, so the diagnosability analysis and diagnoser construction can be performed. We implemented the methodology in a network diagnosis problem.

5. REFERENCES

[1] R. Alur, C. Courcoubetis, N. Halbwachs, T.A. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *Theoretical Computer Science*, 138:3–34, 1995.

[2] S. Bohacek, J.P. Hespanha, J. Lee, and K. Obraczka. A hybrid systems modeling framework for fast and accurate simulation of data communication networks. In *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '03, pages 58–69, 2003.

[3] M.-O. Cordier, P. Dague, F. Levy, J. Montmain, M. Staroswiecki, and L. Trave-Massuyes. Conflicts versus analytical redundancy relations a comparative analysis of the model based diagnosis approach from the artificial intelligence and automatic control perspectives. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 34(5):2163–2177, 2004.

[4] Y. Deng, A. Rajhans, and A.A. Julius. Strong: A trajectory-based verification toolbox for hybrid systems. In *Quantitative Evaluation of Systems*, volume 8054 of *Lecture Notes in Computer Science*, pages 165–168. Springer Berlin Heidelberg, 2013.

[5] M.D. Di Benedetto, S. Di Gennaro, and A. D'Innocenzo. Diagnosability verification for hybrid automata and durational graphs. In *Decision and Control, 2007 46th IEEE Conference on*, pages 1789–1794, 2007.

[6] M.D. Di Benedetto, S. Di Gennaro, and A. D'Innocenzo. Diagnosability of hybrid automata with measurement uncertainty. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 1042–1047, 2008.

[7] M.D. Di Benedetto, S. Di Gennaro, and A. D'Innocenzo. Discrete state observability of hybrid systems. *Int. J. Robust Nonlinear Control*, 19:1564–1580, 2009.

[8] M.D. Di Benedetto, S. Di Gennaro, and A. D'Innocenzo. Verification of hybrid automata diagnosability by abstraction. *Automatic Control, IEEE Transactions on*, 56(9):2050–2061, Sept 2011.

[9] J.J. Gertler. Survey of model-based failure detection and isolation in complex plants. *Control Systems Magazine, IEEE*, 8(6):3–11, 1988.

[10] M. Heymann, F. Lin, G. Meyer, and S. Resmerita. Analysis of zeno behaviors in hybrid systems. In *In: Proceedings of the 41st IEEE Conference on Decision and Control, Las Vagas, NV (2002)*, pages 2379–2384, 2002.

[11] A.A. Julius, G.E. Fainekos, M. Anand, I. Lee, and G.J. Pappas. Robust test generation and coverage for hybrid systems. In *In Proc. of the 10th International Workshop on Hybrid Systems: Computation and Control*, pages 329–342. Springer, 2007.

[12] W. Liu and I. Hwang. Robust estimation and fault detection and isolation algorithms for stochastic linear hybrid systems with unknown fault input. *Control Theory Applications, IET*, 5(12):1353–1368, August 2011.

[13] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *Automatic Control, IEEE Transactions on*, 40(9):1555–1575, 1995.

[14] D. Thorsley. Diagnosability of stochastic chemical kinetic systems: a discrete event systems approach. In *American Control Conference (ACC), 2010*, pages 2623–2630, June 2010.

[15] D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *Automatic Control, IEEE Transactions on*, 50(4):476–492, April 2005.

[16] S. Tripakis. Fault diagnosis for timed automata. In *Formal Techniques in Real-Time and Fault-Tolerant Systems*, volume 2469 of *Lecture Notes in Computer Science*, pages 205–221. Springer Berlin Heidelberg, 2002.

[17] F. Zhao, X. Koutsoukos, H. Haussecker, J. Reich, and P. Cheung. Monitoring and fault diagnosis of hybrid systems. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 35(6):1225–1240, Dec 2005.